

A Note on (Parallel) Depth- and Breadth-First Search by Arc Elimination

Jesper Larsson Träff
 Vienna University of Technology
 Faculty of Informatics, Institute of Information Systems
 Research Group Parallel Computing
 Favoritenstrasse 16/184-5, 1040 Vienna, Austria
 email: `traff@par.tuwien.ac.at`

May 7, 2013

Abstract

This note recapitulates an algorithmic observation for ordered Depth-First Search (DFS) in directed graphs that immediately leads to a parallel algorithm with linear speed-up for a range of processors for non-sparse graphs. The note extends the approach to ordered Breadth-First Search (BFS). With p processors, both DFS and BFS algorithms run in $O(m/p + n)$ time steps on a shared-memory parallel machine allowing concurrent reading of locations, e.g., a CREW PRAM, and have linear speed-up for $p \leq m/n$. Both algorithms need n synchronization steps.

1 Introduction

Depth- and Breadth-First Search are elementary graph traversal procedures with simple, sequential algorithms [2, 6, 7]. Both procedures pose problems for parallel implementation: the ordered Depth-First Search (DFS) problem is P -complete [4], and therefore unlikely to admit polylogarithmically fast, parallel algorithm using only polynomial resources, whereas for Breadth-First Search (BFS), no work-optimal, polylogarithmically fast parallel algorithm is known. This note re-presents the simple, work-optimal, linear time, parallel algorithm for Depth-First Search by Varman and Doshi [8] (also Vishkin, personal communication), that can give linear speed-up for graphs that are not too sparse, and extends the basic observation to Breadth-First Search where an algorithm with similar properties is given. The idea is simple: instead of examining arcs in the “forwards” direction, as in standard, textbook formulations of DFS and BFS [2], incoming, “backwards” arcs are used to eliminate arcs that are no longer relevant for the search. Whereas the standard algorithms have either conflicts (BFS) and/or dependencies (DFS) that hamper parallelization, arc elimination can be performed fully in parallel.

Let $G = (V, E)$ be a directed graph with $n = |V|$ vertices and $m = |E|$ arcs (directed edges). Vertices are assumed to be numbered consecutively, such that $V = \{0, \dots, n-1\}$. Arcs are ordered pairs of vertices with $\langle u, v \rangle$, $u, v \in V$ denoting the arc directed from u to v .

It will be assumed that the input graph G is given as an n -element array ADJ of adjacency arrays. For each vertex $u \in V$, $\text{ADJ}[u].\text{outdeg}$ stores the out-degree of u , and the target vertex v_i of the i th arc $\langle u, v_i \rangle$ for $0 \leq i < \text{ADJ}[u].\text{outdeg}$ is stored in $\text{ADJ}[u].\text{out}[i]$.

Depth- and Breadth-First Search are procedures for graph traversal starting from a given a start vertex $s \in V$. Both procedures assign traversal numbers to the vertices, indicating the

order in which they are reached. Breadth-First Search additionally computes for each vertex its distance (shortest path in number of traversed arcs) from the start vertex. Both procedures also compute the search tree, which will be represented by a parent pointer. For each vertex $u \in V$, these computed values will be stored in $\text{ADJ}[u].\text{traversal}$, $\text{ADJ}[u].\text{distance}$ and $\text{ADJ}[u].\text{parent}$ (for $u \neq s$), respectively.

The search procedures will modify the input graph by eliminating arcs, and both maintain the following invariant.

Invariant 1 *A vertex v is called visited when it has been assigned its (Depth- or Breadth-First Search) traversal number. Each visited vertex $v \in V$ will have no incoming arcs, that is, there will be no arc $\langle u, v \rangle$ for any $u \in V$.*

In order to maintain Invariant 1 the procedures eliminate incoming arcs when a vertex is being visited. To do this efficiently, each vertex $v \in V$ needs an array storing the vertices $u \in V$ for which there is an arc $\langle u, v \rangle$, as well as the index i of v in the array $\text{ADJ}[u].\text{out}$ such that $v = \text{ADJ}[u].\text{out}[i]$. The arrays $\text{ADJ}[v].\text{in}$ for each $v \in V$ shall store the pairs (u, i) representing the incoming arcs of v in this fashion.

To eliminate the arc $\langle u, v \rangle$ from the adjacency array of u , links (indices) to next and previous non-removed vertices in the array are maintained, imposing a doubly linked list on each of the adjacency arrays. The operation $\text{eliminate}(G, u, i)$ removes the i th vertex in $\text{ADJ}[u].\text{out}$ by linking it out of the doubly linked list. The adjacency array itself is not changed. Next and previous indices are maintained in the arrays $\text{ADJ}[u].\text{next}$ and $\text{ADJ}[u].\text{prev}$; $\text{ADJ}[u].\text{first}$ shall index the first non-eliminated vertex in $\text{ADJ}[u].\text{out}$.

Algorithm 1 shows how to compute the array of incoming arcs and the pointers for the doubly linked adjacency lists. A **for**-construct indicates sequential execution for all values in some index set (in some order), whereas the **par**-construct indicates that the computations for each element in the index set can be performed in parallel by the available processors. All processors are assumed to have access to the same memory, and concurrent reading is allowed. Synchronization is implied at the end of each **par**-construct.

Algorithm 1 Computing incoming arcs for each $v \in V$ and doubly linked adjacency lists.

```

par  $u \in V$  do
   $\text{ADJ}[u].\text{indeg} \leftarrow 0$ ,
   $\text{ADJ}[u].\text{first} \leftarrow 0$ ,
end par
for  $u \in V$  do
  par  $0 \leq i < \text{ADJ}[u].\text{outdeg}$  do
     $v \leftarrow \text{ADJ}[u].\text{out}[i]$ 
     $d \leftarrow \text{ADJ}[v].\text{indeg}$ 
     $\text{ADJ}[v].\text{in}[d] \leftarrow (u, i)$  {Add incoming arc  $\langle u, v \rangle$  to  $v$ }
     $\text{ADJ}[v].\text{indeg} \leftarrow d + 1$ 
     $\text{ADJ}[u].\text{next}[i] \leftarrow i + 1$ 
     $\text{ADJ}[u].\text{prev}[i] \leftarrow i - 1$ 
  end par
end for

```

Lemma 1 *Algorithm 1 computes the array of (u, i) vertex-index pairs representing the incoming arcs for all vertices $v \in V$. It also initializes the doubly linked lists over the adjacency arrays $\text{ADJ}[u].\text{out}$. The algorithm runs in $O(m/p + n)$ time steps with p processors using $O(m + n)$ additional space.*

Proof: In each sequential iteration over the set of vertices, incoming arcs are added to *different* target vertices. For each $u \in V$ this can therefore be done in parallel by the p available processors in $O(d(u)/p)$ time steps, where $d(u)$ is the outdegree of vertex u , provided that all processors can read the start address of the array. The total time is $O(n + \sum_{u \in V} d(u)/p) = O(m/p + n)$. \square

The $\text{ADJ}[u].\text{first}$ indices for each $u \in V$ will be maintained such that $\text{ADJ}[u].\text{first} < \text{ADJ}[u].\text{outdeg}$ indicates a non-empty list of non-eliminated arcs out of u . The `eliminate` operation is straightforward.

2 Depth-First Search

We can now present the parallel Depth-First Search algorithm. The DFS procedure is called with a start vertex $s \in V$ and an initial DFS number a , and computes a DFS tree with reachable vertices numbered successively in DFS order starting from a . Each recursive call visits a new vertex, assigns it a DFS number, establishes Invariant 1 by eliminating, in parallel, all arcs into the vertex, and then recursively DFS numbers the subtree from the first non-eliminated arc $\langle s, v \rangle$ out of s ; outgoing arcs are always considered in the fixed order as given in the adjacency array representation of G . The recursion traverses the vertices in G reachable from s in DFS order. The algorithm is given in detail as Algorithm 2, and is essentially as described by Varman and Doshi [8].

Algorithm 2 Recursive, parallel Depth-First Search from start vertex $s \in V$. Vertices that are reachable from s will be assigned successive DFS numbers starting from a .

```

Procedure DFS( $s, G, a$ ):
  par  $0 \leq i < \text{ADJ}[s].\text{indeg}$  do
     $(u, j) \leftarrow \text{ADJ}[s].\text{in}[i]$ 
    eliminate( $G, u, j$ )
  end par
   $\text{ADJ}[s].\text{traversal} \leftarrow a$  {Vertex  $s$  now visited}
   $a \leftarrow a + 1$ 
  while  $\text{ADJ}[s].\text{first} < \text{ADJ}[s].\text{outdeg}$  do {As long as there are un-eliminated arcs}
     $i \leftarrow \text{ADJ}[s].\text{first}$ 
     $v \leftarrow \text{ADJ}[s].\text{out}[i]$ 
     $\text{ADJ}[v].\text{parent} \leftarrow s$ 
     $a \leftarrow \text{DFS}(v, G, a)$ 
  end while
  return  $a$ 

```

Proposition 1 *Algorithm 2 computes an ordered Depth-First Search numbering and tree in $O(m/p + n)$ time steps using p processors.*

Proof: By Invariant 1 once a vertex v is visited, it will never be considered again, since all arcs into v will have been eliminated. Therefore, each vertex in G that is reachable from s will be visited once. The time complexity is immediate: when a vertex is visited the incoming arcs are eliminated in parallel. Since $\text{ADJ}[u].\text{first}$ will for each vertex be the index of the first adjacent vertex v where the arc $\langle u, v \rangle$ has not been eliminated, the order in which vertices are visited is the same as standard DFS search procedures, from which the correctness follows. \square

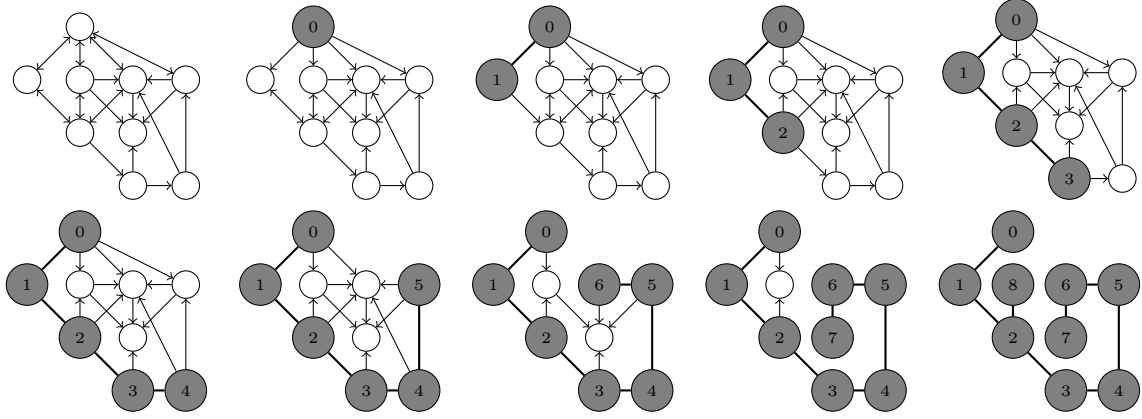


Figure 1: A sample graph $G = (V, E)$ and the DFS traversal as per Algorithm 2 starting from the topmost node. Arcs are examined in counter-clockwise order, starting from lower left. Node labels are the DFS numbers, and tree edges are indicated as heavy, undirected edges. Arcs disappear as they are being eliminated, leaving at the end the heavy DFS tree.

The algorithm also computes a DFS tree by setting parent pointers for the visited vertices. Note that it can easily be extended to classify arcs into backwards, forwards, tree and cross arcs, as sometimes desirable by a DFS traversal, without changing the time bounds. An example execution of the algorithm is given in Figure 1.

3 Breadth-First Search

The arc elimination idea can also be used for parallel Breadth-First Search, as shown in Algorithm 3. The algorithm has the same structure as standard, “forwards” BFS [2], but performs parallel arc elimination as new, unexplored vertices are added to the queue for the next level. This ensures that each reachable vertex is explored once.

Proposition 2 *Algorithm 3 computes an ordered Breadth-First Search numbering and tree in $O(m/p + n)$ time steps using p processors.*

Proof: Again by Invariant 1, once a vertex has been visited it will never be considered again, and therefore arc elimination is performed once for each reachable vertex. From this the time bound follows. For each vertex in Q for some level, the un-eliminated arcs are considered in order determined by the representation of G , and as each unvisited vertex is put into the queue Q' for the next level, all incoming arc are eliminated. This in particular ensures that there are no arcs between vertices in Q' . As for standard BFS, all vertices in Q before the start of an iteration of the innermost repeat loop have the same distance to the source vertex, from which correctness follows. \square

It is especially worth noticing that there are no arcs between nodes in Q' , the queue being filled for the next iteration. An example execution of the algorithm is given in Figure 2. The important property of BFS is that vertices are explored in least recently visited order; it is therefore, also in the arc elimination algorithm, possible to dispense with the explicit next level queue Q' and do with only a single **repeat**-loop [7].

Algorithm 3 Parallel Breadth-First Search from start vertex $s \in Q$. Vertices that are reachable from s will be assigned a BFS number starting from a ; also the distance from s (in smallest number of arcs) will be computed.

```

Procedure BFS( $s, G, a$ ):
  par  $0 \leq i < \text{ADJ}[s].\text{indeg}$  do
     $(u, j) \leftarrow \text{ADJ}[s].\text{in}[i]$ 
     $\text{eliminate}(G, u, j)$ 
  end par
   $l \leftarrow 0$ 
   $\text{ADJ}[s].\text{traversal} \leftarrow a$ 
   $\text{ADJ}[u].\text{distance} \leftarrow l$ 
   $Q.\text{enqueue}(s)$  {Start vertex  $s$  visited}
   $Q' \leftarrow \emptyset$ 
  repeat
     $l \leftarrow l + 1$  {Next level}
    repeat
       $u \leftarrow Q.\text{deque}()$ 
      while  $\text{ADJ}[u].\text{first} < \text{ADJ}[u].\text{outdeg}$  do {As long as there are un-eliminated arcs}
         $i \leftarrow \text{ADJ}[u].\text{first}$ 
         $v \leftarrow \text{ADJ}[u].\text{out}[i]$ 
        par  $0 \leq j < \text{ADJ}[v].\text{indeg}$  do
           $(w, k) \leftarrow \text{ADJ}[v].\text{in}[j]$ 
           $\text{eliminate}(G, w, k)$ 
        end par
         $a \leftarrow a + 1$  {Next vertex}
         $\text{ADJ}[v].\text{traversal} \leftarrow a$ 
         $\text{ADJ}[v].\text{distance} \leftarrow l$ 
         $\text{ADJ}[v].\text{parent} \leftarrow u$ 
         $Q'.\text{enqueue}(v)$  {Vertex  $v$  has now been visited, enqueue for next level}
      end while
    until  $Q = \emptyset$ 
     $Q \leftarrow Q'$ 
  until  $Q = \emptyset$ 

```

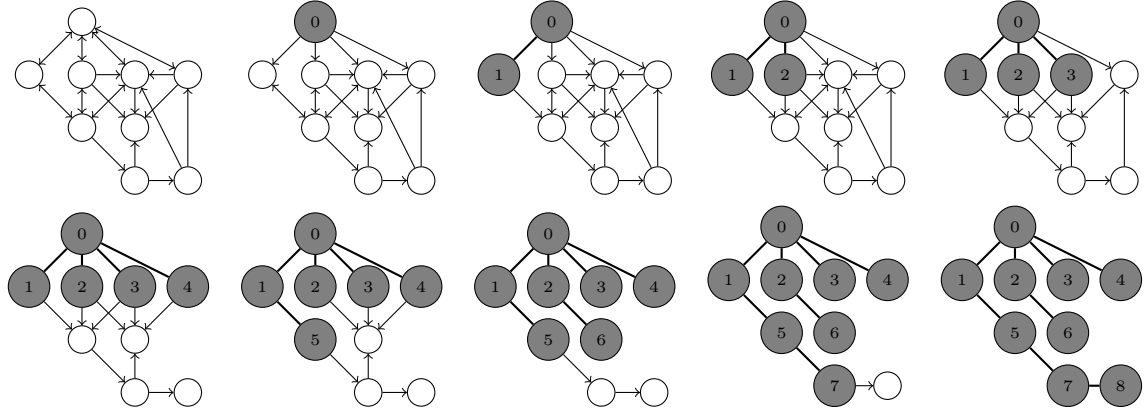


Figure 2: The sample graph $G = (V, E)$ and the BFS traversal as per Algorithm 3 starting from the topmost node. Arcs are examined in counter-clockwise order, starting from lower left. Node labels are the computed BFS numbers, and tree edges are indicated as heavy, undirected edges. Arcs disappear as they are being eliminated, leaving at the end the heavy BFS tree.

4 Discussion

The time bounds for both Depth- and Breadth-First Search algorithms guarantee linear speed-up when $p \leq m/n$, or equivalently $m \geq pn$; that is, good speed-up is possible for graphs with average degree larger than the number of processors. The algorithms presented here are complementary to the standard, textbook, “forwards” procedures for DFS and BFS [2]. Standard DFS where arcs are examined only in the forwards direction has no parallelism; in contrast, the algorithm given here can perform the arc elimination fully in parallel. Typical, parallel Breadth-First Search algorithms exploit parallelism mostly by considering active vertices in the queue for each level in parallel, see, e.g., [1, 3, 5]. Although the forward edges also be can be explored in parallel, compaction or other data structure operations are necessary for resolving/avoiding update conflicts and maintaining the queue for the next level. The parallel running time of such algorithms will typically be bounded by the diameter of the graph, but either at the cost of more work incurred by data structure operations, or by requiring stronger, atomic operations. The arc elimination approach does not require either of these means (data structures, compaction, atomic operations), and has exploitable parallelism independent of the BFS structure of the graphs, as long as the total number of edges m satisfies $m \geq np$.

References

- [1] D. A. Bader and K. Madduri. Designing multithreaded algorithms for breadth-first search and st-connectivity on the Cray MTA-2. In *International Conference on Parallel Processing (ICPP)*, pages 523–530, 2006.
- [2] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms*. MIT Press, third edition, 2009.
- [3] C. E. Leiserson and T. B. Schardl. A work-efficient parallel breadth-first search algorithm (or how to cope with the nondeterminism of reducers). In *22nd Annual ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 303–314, 2010.

- [4] J. H. Reif. Depth-first search is inherently sequential. *Information Processing Letters*, 20:229–234, 1985.
- [5] J. Shun and G. E. Blelloch. Ligra: a lightweight graph processing framework for shared memory. In *ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP)*, pages 135–146, 2013.
- [6] R. E. Tarjan. Depth-first search and linear graph algorithms. *SIAM Journal on Computing*, 1(2):146–160, 1972.
- [7] R. E. Tarjan. *Data Structures and Network Algorithms*. Society of Industrial and Applied Mathematics (SIAM), 1983.
- [8] P. J. Varman and K. Doshi. Improved parallel algorithms for the depth-first search and monotone circuit value problems. In *15th ACM Conference on Computer Science*, pages 175–182, 1987.